Locust Olfaction

Synchronous Oscillations in Excitatory and Inhibitory Groups of Spiking Neurons

David C. Sterratt*

Institute for Adaptive and Neural Computation, Division of Informatics,
University of Edinburgh,
dcs@cogsci.ed.ac.uk,
http://www.cogsci.ed.ac.uk/~dcs

Abstract. During odour recognition, excitatory and inhibitory groups of neurons in the second stage of the locust olfactory system, the antennal lobe (AL), fire alternately. There is little spread in the firing times within each group. Locust anatomy and physiology help to pin down all parameters apart from the weights in a coarse spiking neuron model of the AL. The time period and phase of the group oscillations do however constrain the weights; this paper investigates how.
I generalise the spiking neuron locking theorem [3] to derive conditions that allow stable synchronous firing within multiple groups. I then apply the general result to the AL model. The most important result is that for a general form of postsynaptic potential (PSP) function the excitatory and inhibitory neuronal populations cannot fire alternately at certain time periods and phases, regardless of the size of the weights between and within groups.

The dynamics of groups of spiking neurons connected by purely excitatory or purely inhibitory synapses and pulsatile connections have been well studied (for example [3,7,8,9]). The Wilson-Cowan oscillator [11] has inspired a number of models of the dynamics of the mass activity of reciprocally connected groups of inhibitory and excitatory neurons. However, the dynamics of reciprocally connected groups of excitatory and inhibitory spiking neurons do not appear to have been studied, perhaps due to lack of biological motivation.

A system that provides some motivation is the peripheral olfactory system of the locust *Schistocerca americana*, which has been studied by Laurent and co-workers [4]. Its anatomy and mass activity is summarised in Fig. 1. Each antenna contains about 50 000 excitatory primary olfactory sensory neurons, called receptor neurons (RNs), which project into the ipsilateral AL. The AL contains ∼830 excitatory projection neurons (PNs) and ∼300 inhibitory local neurons (LNs). There appear to be connections from RNs to PNs and LNs and reciprocally within and between the PN and LN groups. The connections are

* This work was carried out under the supervision of David Willshaw and Bruce Graham, with the financial support of the Caledonian Research Foundation. I would like to thank Leslie Smith for comments on an earlier draft.

S. Wermter et al. (Eds.): Emergent Neural Computational Architectures, LNAI 2036, pp. 270–284, 2001.

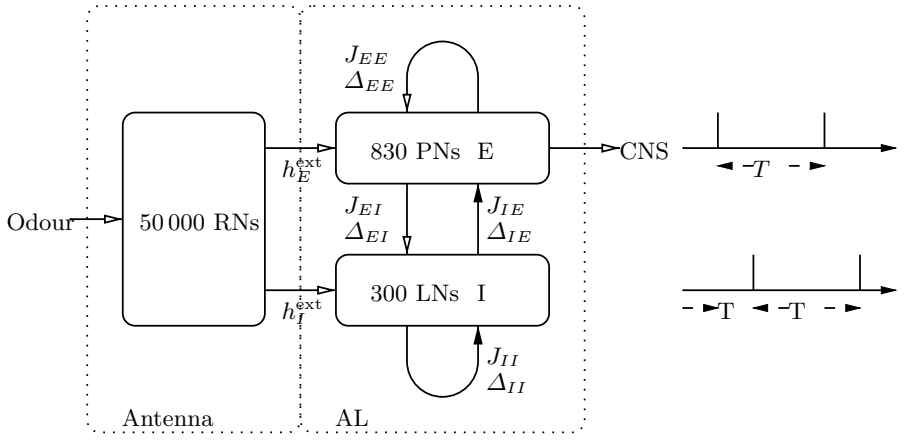

Fig. 1. The basic anatomy and activity of the locust AL. The left hand part of the figure shows the anatomy as described in the text and the symbols used later on in the paper. Open arrowheads denote excitatory connections and filled ones indicate inhibitory ones. The right hand part shows the group activity of the PNs (top) and LNs (bottom).

organised in 1000 groups of synapses called glomeruli, and the mean number of each type of neuron projecting into a glomerulus is known.

When no odour reaches the antenna, the RNs fire spontaneously at a low rate. When odour is present, their firing rate either increases, decreases or remains the same, depending on the identity of the RN and the odour. The PNs are conventionally spiking neurons, whereas the LNs are axon-less neurons that produce graded action potentials. With odour present at the antenna, subsets of the PNs and LNs fire; otherwise they are mainly quiescent. The overall activity of the both sets of neurons oscillates at about 20 Hz and the excitatory PNs lead the inhibitory LNs by quarter of a cycle [5]. This frequency and phase seem to be fairly independent of the concentration of the odour. Furthermore, it appears that there is usually only one spike per PN per cycle [6].

Experimental evidence suggests that the spatiotemporal pattern of activation of the PNs over a period of about 1 s encodes the identity of odours presented to the system [4,5]. The oscillations seem to be important for refining the representation of the odour, as when they are pharmaceutically abolished honeybees find it harder to discriminate between similar odours.

In order to understand the mechanism of the oscillations and spatiotemporal patterns I have been modelling the RNs and AL. Although the LNs are not true spiking neurons, one way of modelling them is as spiking neurons since it appears that their effect is similar to inhibitory postsynaptic potentials (IPSPs) from synchronously-firing inhibitory neurons [5]. As there is not enough data to produce a biophysical model of each type of neuron, I used the spike response model (SRM) [2] and fitted the PSPs, refractory curves, axonal delays

and thresholds to experimental data. From the glomerular structure I estimated the connectivity, but there is no data that directly constrains the weights.

However, the (mean) weights clearly affect the time period and phase of oscillation. The first question this paper sets out to answer is how we can use time period and phase of oscillation to reduce the degrees of freedom in the model by constraining the weights.

To answer this question we need to solve a generalised version of the locking problem [3]. The problem concerns a network comprising homogeneous groups of neurons. They are homogeneous in the sense that the parameters of all neurons in a group are the same and there is a certain kind of homogeneity in the connections between different groups, which will be defined more precisely later. What conditions are required to achieve a network state where the neurons of each group fire synchronously with one another and each group fires in turn? The problem has been solved for one infinitely-large group [3] using SRM neurons.

A second question is how the time period and phase of oscillation depend on the input strength to the system.

In Sect. 1, I briefly introduce the SRM. In Sect. 2, I give the mathematical form of the constraints for multiple groups and confirm that the locking theorem can be extended. In Sect. 3, I present the implications of the constraints for the excitatory-inhibitory oscillator. In Sect. 4, a network simulation with one set of parameters is consistent with most of the analytical and numerical predictions but shows some differences, possibly due to the simulation method. Finally, in Sect. 5, I discuss the conclusions and suggest some further work.

1 The Spike Response Model

The SRM [2] provides a framework for describing and analysing networks of spiking neurons. It comes in two forms, both of which can be derived by integrating the integrate-and-fire (IF) neuron equations. I use the simpler SRM model as it can describe adaptation.

In a network of N neurons, the membrane potential $h_i(t)$ at time t of each neuron i is given by

$$h_i(t) = h_i^{\text{syn}}(t) + h_i^{\text{ref}}(t) + h_i^{\text{ext}}(t) \tag{1}$$

where $h_i^{\text{syn}}(t)$ is the *postsynaptic* potential of neuron i, $h_i^{\text{ref}}(t)$ is the *refractory* potential and $h_i^{\text{ext}}(t)$ is the *external* (input) potential. When the membrane potential reaches a threshold θ_i the neuron fires.

The postsynaptic potential comprises PSPs caused by spikes from the neurons in the network:

$$\sum_{i=1}^{N}\sum_{f=1}^{F} h_i^{\text{syn}}(t) = J_{ij}\epsilon_{ij}\left(t - t_j^{(f)}\right) \quad , \tag{2}$$

where J_{ij} is the weight from neuron j to neuron i, $\epsilon_{ij}(s)$ is the PSP function from neuron j to neuron i, $t_i^{(f)}$ represents the fth last firing time of neuron i and

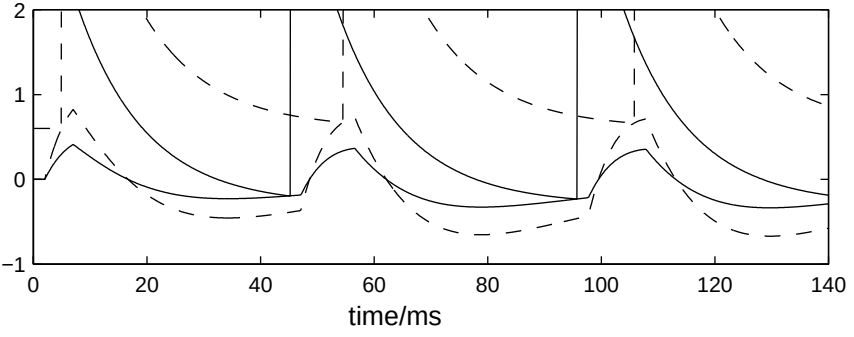

Fig. 2. An example of the synaptic potentials (lower traces) and effective threshold $\theta - h^{\mathrm{ext}}(t) - h^{\mathrm{ref}}(t)$ (upper traces) of neurons in the excitatory (solid line) and inhibitory (dashed line) groups. Each neuron fires when its membrane potential crosses the effective threshold. The parameters are: $\tau_{\mathrm{E}} = 10\,\mathrm{ms}$, $\tau_{\mathrm{I}} = 15\,\mathrm{ms}$, $\Delta = 2\,\mathrm{ms}$, $J_{\mathrm{EE}} = 0.5$, $J_{\mathrm{EI}} = 0.5$, $J_{\mathrm{IE}} = 1$, $J_{\mathrm{II}} = 1$, $F = 3$, $\theta = 0$, $h_{\mathrm{E}}^{\mathrm{ext}} = 0.3$ and $h_{\mathrm{I}}^{\mathrm{ext}} = -0.6$. The resulting time period and phase are $T = 51.3\,\mathrm{ms}$ and $\phi = 0.200$. For the meaning of the parameters see Sect. 3.

F is the number of spikes remembered. A common form for the PSP function is the alpha function $\epsilon_{ij}(s) = (s - \Delta_{ij})/\tau e^{-(t-\Delta_{ij})/\tau}$ where Δ_{ij} is the delay from neuron j to neuron i.

The refractory potential depends on the firing time of the neuron:

$$h_i^{\mathrm{ref}}(t) = \sum_{f=1}^{F} \eta_i\Big(t - t_i^{(f)}\Big) \quad . \tag{3}$$

Here the refractory function $\eta(s)$ derives from the reset of an IF neuron without an absolute refractory period, but can be used to simulate absolute refractory periods and after-potentials. When derived from the IF model its derivative is always positive ($d\eta_i/ds > 0$). This property is important for the locking properties later on, and is called the *standard dynamics* [3]. In this paper I will use an exponential function derived from recordings from locust AL PNs [10]:

$$\eta(s) = \begin{cases} 0 & s \le 0 \\ \exp(1.5 - s/12) & s > 0 \end{cases} \quad . \tag{4}$$

Figure 2 shows an example of the synaptic potential and effective threshold of an excitatory and inhibitory neuron in a model of the AL.

2 Constraints on Stable Synchronous Oscillations

Suppose we have a network comprising N neurons organised into M groups defined by the connectivity, delays, thresholds, and PSP and refractory functions.

I denote the set of neurons in the lth group by $\mathcal{G}_l$. The connectivity is such that the total presynaptic connection strength from each group is the same for all neurons in a group. We can express this mathematically by

$$\sum_{j \in \mathcal{G}_m} J_{ij} = \tilde{J}_{lm} \ , \tag{5}$$

where $\tilde{J}_{lm}$ denotes the total weight from the neurons of group m onto a neuron in group l. We further demand that the PSP functions and the weights are the same ($\epsilon_{ij}(s) = \tilde{\epsilon}_{lm}(s)$ and $\Delta_{ij} = \tilde{\Delta}_{lm}$) for all neurons $i \in \mathcal{G}_l$ and $j \in \mathcal{G}_m$. Finally, we require that the thresholds and refractory functions for neurons in a group be equal: $\theta_i = \tilde{\theta}_l$ and $\eta_i(s) = \tilde{\eta}_l(s)$ for $i \in \mathcal{G}_l$.

In a multi-group locked state the neurons of each group fire synchronously and periodically with a time period T. We assume that group l fires at $(k+\phi_l)T$ for integer k where ϕ_l is a phase variable in the range [0,1). It is convenient to arrange the phases in ascending order and set $\phi_1 = 0$ so that $0 = \phi_1 < \phi_2 < \ldots < \phi_M$.

Now assume that the $l-1$th group has fired and that we are waiting for the lth group to fire (or that the Mth group has fired and we are waiting for the first group to fire). Although time is arbitrary, we will assume that we are in the first cycle after $t = 0$, that is that $\phi_{l-1} < t < \phi_l$ (or $\phi_l < t < T$ if we are waiting for the first group). For any neuron $i \in \mathcal{G}_l$ the expected synaptic and refractory potentials are

$$h_i^{\mathrm{syn}}(t) = \sum_{l=1}^{m-1} \tilde{J}_{lm} \sum_{k=0}^{1-F} \tilde{\epsilon}_{lm}(t-(k+\phi_m)T) + \sum_{l=m}^{M} \tilde{J}_{lm} \sum_{k=-1}^{-F} \tilde{\epsilon}_{lm}(t-(k+\phi_m)T) \tag{6}$$

and

$$h_i^{\mathrm{ref}}(t) = \sum_{k=-1}^{-F} \tilde{\eta}_m(t-(k+\phi_l)T) \ . \tag{7}$$

Group l should next fire when the membrane potential reaches the threshold θ_l. If the threshold has been set correctly this will be at $t = \phi_l T$, leading to the self-consistency or *threshold* condition:

$$\tilde{\theta}_l = h_i(\phi_l T) \ . \tag{8}$$

A more precise requirement for the existence of coherent solutions is that for all $i \in \mathcal{G}_l$,

$$\phi_l T = \inf(t > (\phi_{l-1})T | h_i(t) = \tilde{\theta}_l) \ . \tag{9}$$

As $h_i(t)$ reaches $\tilde{\theta}_l$ from below, this implies that a necessary condition for stability is

$$\dot{h}_i(\phi_l T) > 0 \ . \tag{10}$$

I refer to this as the *rising potential condition*.

Although these conditions show that synchronised solutions exist, they do not show that they are stable. In order to show this, we have to show that a

small perturbation to the firing time of a neuron from the group firing in a cycle time will decrease in the next cycle. This linear perturbation analysis applied to a homogeneous network of neurons with standard dynamics yields the *locking theorem* [3]:

Theorem 1. *In a spatially homogeneous network of spiking neurons with standard dynamics, a necessary and, in the limit of a large number n of presynaptic neurons ($n \to \infty$), also sufficient condition for a coherent oscillation to be asymptotically stable is that firing occurs when the postsynaptic potential arising from all previous spikes is increasing in time.*

In a system with standard dynamics this is a stricter condition than the rising potential condition since the positive refractory derivative contributes to making the derivative of the membrane potential positive.

This result does not apply to more than one group of simultaneously firing neurons. However, it is straightforward to extend the proof to prove part of the the *infinite-limit extended locking theorem*:

Theorem 2. *In a network of spiking neurons with standard dynamics with M homogeneous groups of neurons such that all neurons in group l receive the same total postsynaptic weight from a large number of presynaptic neurons N_m in each group m a necessary and sufficient condition for constant phase, periodic, coherent oscillation to be asymptotically stable is that firing occurs when the postsynaptic potential arising from all previous spikes is increasing in time.*

Appendix A contains the proof. The mathematical expression of the extended locking theorem is that

$$\dot{h}_i^{\text{syn}}(\phi_l T) > 0 \quad \text{for } i \in \mathcal{G}_l \ . \tag{11}$$

3 Applying the Constraints to Excitatory and Inhibitory Groups

It is now a straightforward task to apply the extended locking theorem to one excitatory and one inhibitory group of neurons, a system that approximates the AL.

To make the notation clearer, I will refer to the excitatory group as E and the inhibitory one as I. Later it will be convenient to work with positive weights, so I refer to the group connection weights which should be called $\tilde{J}_{\text{EE}}$, $\tilde{J}_{\text{EI}}$ and so on as $J_{\text{EE}}, -J_{\text{EI}}, J_{\text{IE}}$ and $-J_{\text{II}}$. (I drop the tildes from here on as there is no further need to refer to the weights between individual neurons.) I take the delays to be uniformly Δ. I assume that PSP functions do not depend on the postsynaptic neurons and denote the excitatory postsynaptic potential (EPSP) function by $\epsilon_{\text{E}}(t)$ and the IPSP function by $\epsilon_{\text{I}}(t)$. I will assume that they are zero until $t = \Delta$, rise monotonically to a maximum at time $t = \tau_{\text{E}}$ or $t = \tau_{\text{I}}$ and then fall monotonically back to zero. The refractory functions, denoted $\eta_{\text{E}}(t)$ and $\eta_{\text{I}}(t)$, both have positive derivatives so the network has standard dynamics.

To simplify the analysis I will assume that neurons only remember their last firing times, that is $F = 1$. This is reasonable if the membrane potential due to

a second spike is much smaller than the membrane potential due to the first, which is true when the time period is long compared to the PSP time constant. In any case, it is straightforward to generalise to longer spike memories. I will also assume that the refractory function complies with the formal version of the threshold condition (9). This means the absolute refractory period is less than the desired period T and that it prevents multiple firing during a cycle.

Since with the standard dynamics the stability condition is stronger than the rising potential condition, we can ignore the rising potential condition. Of the two remaining conditions, the stability condition is the more fundamental, since it constrains only the weights, whereas the threshold condition also constrains the input levels. I therefore explore the consequences of the stability condition first, in Sect. 3.1, before moving onto the consequences of the threshold condition in Sect. 3.2.

3.1 Consequences of the Stability Constraint

In the alternating state, network activity oscillates with a time period of T and the excitatory group leads the inhibitory group by ϕT, where ϕ is a phase variable with the range [0,1). Applying the stability constraint (11) yields

$$J_{\mathrm{EE}}\dot{\epsilon}_{\mathrm{E}}(T) - J_{\mathrm{EI}}\dot{\epsilon}_{\mathrm{I}}\left((1-\phi)T\right) > 0 \tag{12}$$

for the excitatory group and

$$J_{\mathrm{IE}}\dot{\epsilon}_{\mathrm{E}}(\phi T) - J_{\mathrm{II}}\dot{\epsilon}_{\mathrm{I}}(T) > 0 \ . \tag{13}$$

for the inhibitory group.

The consequence of these inequalities is that there are combinations of time periods and phase for which there can be no stable firing, regardless of the size of the weights. This is because either inequality is not satisfied when the excitatory derivative is less than or equal to zero and the inhibitory derivative is greater than or equal to zero. For the condition on the excitatory group (12) this happens when $T < \Delta$ or $T \geq \tau_E + \Delta$ and $0 \leq (1-\phi)T \leq \tau_{\mathrm{I}} + \Delta$. The condition on the inhibitory group (13) is violated when $T \leq \tau_{\mathrm{I}} + \Delta$ and $\phi T \geq \tau_{\mathrm{E}} + \Delta$ or $\phi T < \Delta$. As it is enough for one group to become unstable to plunge the whole system into instability, the forbidden combinations of T and ϕ are given by the union of the regions defined by the above inequalities. They can be displayed in a plot of ϕ versus T (referred to as T-ϕ space from now on) as in Fig. 3.

The second consequence of the inequalities (12) and (13) is that for those combinations of T and ϕ for which activity is not inherently unstable, stability depends on appropriate weights. Thus for any possible pair of T and ϕ there will be a region of the four-dimensional weight space that allows stable firing. We can make a major simplification by noting that the inequalities concern the *ratios* $\hat{J}_{\mathrm{EE}} := J_{\mathrm{EE}}/J_{\mathrm{EI}}$ and $\hat{J}_{\mathrm{IE}} := J_{\mathrm{IE}}/J_{\mathrm{II}}$ and finding the regions of a two-dimensional weight-ratio space that allow stable firing. Since the weights are by definition positive, the weight ratios are always greater than or equal to zero.

The stability condition (12) imposes different constraints on T, ϕ and the weight ratio $\hat{J}_{\mathrm{EE}}$ depending on the sign of $\dot{\epsilon}_{\mathrm{E}}(T)$, and hence on the size of T compared to Δ and $\tau_{\mathrm{E}} + \Delta$.

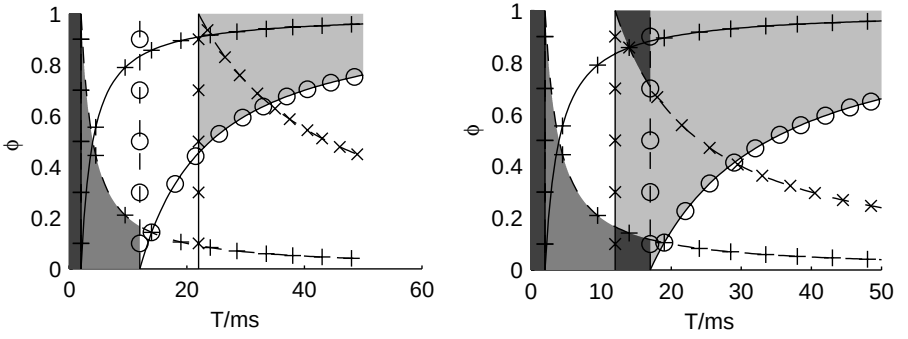

Fig. 3. Forbidden regions of the T-ϕ plane for two sets of parameters. The light regions are unstable due to the excitatory group becoming unstable, medium-grey areas unstable due to the inhibitory group and the darkest areas unstable due to both groups. Dashed lines pertain to the excitatory group, dotted lines to the inhibitory group and dash-dotted lines to both. Plus signs indicate where condition (14) and its inhibitory counterpart become true. Crosses indicate zeros in the denominators and circles zeros in the numerators of equations (15) and (16) and their inhibitory counterparts. Left: $\Delta = 2\,\text{ms}$, $\tau_\text{E} = 20\,\text{ms}$ and $\tau_\text{I} = 10\,\text{ms}$. Right: $\Delta = 2\,\text{ms}$, $\tau_\text{E} = 10\,\text{ms}$ and $\tau_\text{I} = 15\,\text{ms}$ (estimated locust parameters).

1. For $T < \Delta$ or $T = \tau_E + \Delta$, $\dot{\epsilon}_\text{E}(T) = 0$ and from (12) we can deduce that

$$\dot{\epsilon}_\text{I}((1-\phi)T) = 0 \ , \tag{14}$$

   so that stability condition is not satisfied. This is the region $T < \Delta$ where both conditions are violated (see Fig. 3).
2. If $\Delta < T < \tau_\text{E} + \Delta$, $\dot{\epsilon}_\text{E}(T) > 0$ and

$$\hat{J}_\text{EE} > \frac{\dot{\epsilon}_\text{I}((1-\phi)T)}{\dot{\epsilon}_\text{E}(T)} \ . \tag{15}$$

   When $\dot{\epsilon}_\text{I}((1-\phi)T) \leq 0$ (for $(1-\phi)T \leq 0$ or $T \geq \tau_\text{I} + \Delta$) this equation is satisfied for any weight ratio $\hat{J}_\text{EE}$ (and thus any weights J_EE and J_EI) since weight ratios are always positive.
3. When $T > \tau_\text{E} + \Delta$, $\dot{\epsilon}_\text{E}(T) < 0$ and equation (15) does not hold, since we have to reverse the sign of the inequality when rearranging (12). Instead,

$$\hat{J}_\text{EE} < \frac{\dot{\epsilon}_\text{I}((1-\phi)T)}{\dot{\epsilon}_\text{E}(T)} \ . \tag{16}$$

   When $\epsilon_\text{I} \geq 0$ (for $0 < (1-\phi)T \leq \tau_\text{I} + \Delta$) the inequality cannot be satisfied for positive weight ratios. This corresponds to the forbidden region where $T > \tau_\text{E} + \Delta$.

We can derive similar inequalities for the inhibitory group from (13).

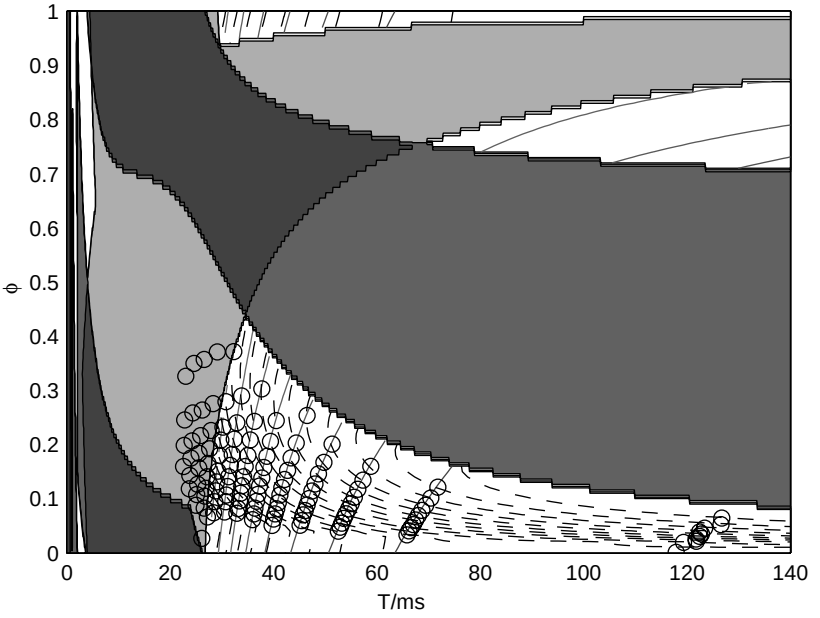

Fig. 4. Unstable regions of T-ϕ space and mapping from $h_{\mathrm{E}}^{\mathrm{ext}}$-$h_{\mathrm{I}}^{\mathrm{ext}}$ space onto T-ϕ space for a pair of weight ratios. The shaded regions are unstable, with the shade representing the groups whose stability conditions are violated: the light region corresponds to the excitatory group, the intermediate region to the inhibitory group, and the darkest region to both groups. The unshaded regions with lines plotted in them are stable given the correct input. The lines represent the levels of external excitation required to produce the values of T and ϕ. Lighter shaded lines correspond to higher values. The solid lines represent constant $h_{\mathrm{E}}^{\mathrm{ext}}$ and the dashed lines represent constant $h_{\mathrm{I}}^{\mathrm{ext}}$. The circles represent simulation values obtained by setting the input levels to those of the lines. They agree well with the theoretical results, although they appear in the theoretically unstable area too. Parameters: $\tau_{\mathrm{E}} = 10\,\mathrm{ms}$, $\tau_{\mathrm{I}} = 15\,\mathrm{ms}$, $\Delta = 2\,\mathrm{ms}$, $J_{\mathrm{EE}} = 0.5$, $J_{\mathrm{EI}} = 0.5$, $J_{\mathrm{IE}} = 1$, $J_{\mathrm{II}} = 1$ and $F = 3$.

The third consequence of the inequalities (12) and (13) complements the second. Any pair of weight ratios limits the stable combinations of T and ϕ even more than the inherent stability. This follows from setting $\hat{J}_{\mathrm{EE}}$ in the inequalities (14), (15) and (16) and solving for T and ϕ. By contrast to the inherently unstable areas, the solution does depend on the exact form of the PSP functions; we cannot simply consider the signs of the PSP gradients.

With alpha function PSPs (defined in Sect. 1) there appears to be no analytical solution to the inequalities (14), (15) and (16) and their inhibitory counterparts. We must solve both sets of inequalities numerically to find the regions each allow. The intersection of these regions is the allowable region. Figure 4 shows the regions of stability and instability for a particular pair of weight ratios.

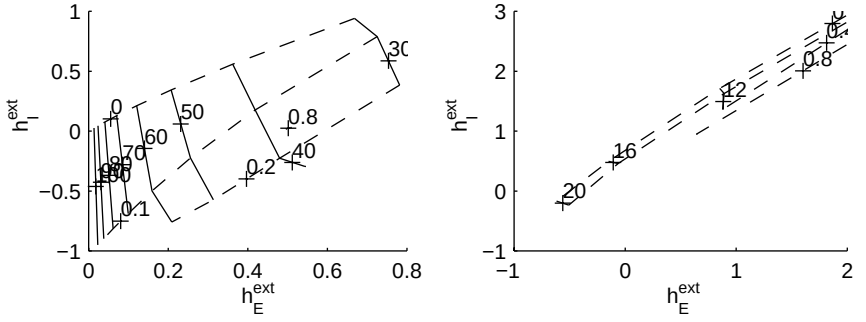

Fig. 5. The mapping from T-ϕ space into $h_{\mathrm{E}}^{\mathrm{ext}}$-$h_{\mathrm{I}}^{\mathrm{ext}}$ space for two sets of parameters. Solid lines correspond to constant T and dashed lines to constant ϕ; numbers between 0 and 1 label some of the constant-T lines with the phase and multiples of 10 label some of the constant=ϕ lines. Left: τ_{E} =10 ms, τ_{I} =15 ms, Δ = 2 ms, $J_{\mathrm{EE}} = 0.5$, $J_{\mathrm{EI}} = 1$, $J_{\mathrm{IE}} = 0.5$, $J_{\mathrm{II}} = 1$, $F = 3$. Right: τ_{E} =20 ms, τ_{I} =10 ms, $\Delta = 3$ ms, $J_{\mathrm{EE}} = 1$, $J_{\mathrm{EI}} = 0.5$, $J_{\mathrm{IE}} = 1$, $J_{\mathrm{II}} = 0.5$, $F = 3$.

3.2 Consequences of the Threshold Constraint

The threshold condition (8) means that

$$\theta_{\mathrm{E}} = J_{\mathrm{EE}}\epsilon_{\mathrm{E}}(T) - J_{\mathrm{EI}}\epsilon_{\mathrm{I}}((1-\phi)T) + \eta_{\mathrm{E}}(T) + h_{\mathrm{E}}^{\mathrm{ext}} \tag{17}$$

holds for the excitatory group and that

$$\theta_{\mathrm{I}} = J_{\mathrm{IE}}\epsilon_{\mathrm{E}}(\phi T) - J_{\mathrm{II}}\epsilon_{\mathrm{I}}(T) + \eta_{\mathrm{I}}(T) + h_{\mathrm{I}}^{\mathrm{ext}} \tag{18}$$

holds for the inhibitory group.

These equations tell us directly how much activation is required to sustain activity in a stable region of T-ϕ space. Figure 5 shows the mapping of a grid in T-ϕ space into $h_{\mathrm{E}}^{\mathrm{ext}}$-$h_{\mathrm{I}}^{\mathrm{ext}}$ space for a particular set of weights and thresholds. There are a few points to notice here.

1. As the excitatory or inhibitory external potential is increased, the time period decreases. It seems intuitive that more excitatory input should speed up the oscillations, but not that more inhibitory input should. The reason is this happens is that more excitatory input makes the inhibitory neurons fire earlier, which in turn release the excitatory neurons from their inhibition earlier.
2. More input to the excitatory neurons means they increase their lead over the inhibitory neurons. We would expect more excitatory input to make the excitatory neurons fire earlier, but without doing calculations we would not know whether it brings forward the inhibitory neurons' firing by more or less. By contrast, more input to the inhibitory neurons reduces the excitatory neurons' lead.

3. The input to both groups can be above or below threshold. It is somewhat surprising that the excitatory group can have below-threshold input as it requires above-threshold external input to start firing. However, once it is firing the self-excitation is sufficient to keep it firing and the external input can be reduced. So below-threshold input to the excitatory group is a hysteresis state.

To find the amount of external input required to obtain a particular time period and phase we must solve (17) and (18) numerically. The lines in Fig. 4 show a grid in $h_{\mathrm{E}}^{\mathrm{ext}}$-$h_{\mathrm{I}}^{\mathrm{ext}}$ space mapped into allowed regions of T-ϕ space. We can see that changes to either $h_{\mathrm{E}}^{\mathrm{ext}}$ or $h_{\mathrm{I}}^{\mathrm{ext}}$ change both T and ϕ.

4 Simulations

To test whether the analysis is correct, I simulated a network with 100 excitatory and 100 inhibitory neurons and the parameters in the caption of Fig. 4 for varying values of external input. I analysed the results to find whether there was stable firing, and if so, the time period and phase of oscillations. In order to make sure that the network did not stay in a stable state by chance, I added some noise to the firing times by making the firing probabilistic [2]. The simulation results are shown in Fig. 4, along with the theoretical results.

In the allowed regions they fit well with the theoretical predictions. However, the network appears to fire stably even when it should not according to the theory. It is not clear why this happens, though it is possible that the noise may stabilise the network in these areas.

5 Discussion

This paper started with a model of the first stage of olfactory processing in the locust, the AL. This model raised the problem of choosing weights in a network of excitatory and inhibitory spiking neurons that produce the right time period and phase of oscillation. In order to do this I had to extend the locking theorem [3] slightly. By applying the extended locking theorem to a network comprising excitatory and inhibitory homogeneous groups I showed:

- that the excitatory and inhibitory neurons can only fire synchronously and alternately for certain time periods and phases of oscillation;
- how to find weights that produce a certain, allowed, time period and phase of oscillation;
- that oscillations are possible for below-threshold input to inhibitory neurons and, when the network is in a hysteresis state, for below-threshold input to excitatory neurons; and
- how the time period and phase depend on the input and vice versa.

Simulations confirm most of these results, but, puzzlingly, seem to be stable when the theory predicts they should not. This may be due to the symmetry-breaking noise included in the simulations.

I am not aware of any model of the locust AL at a similar level to mine, nor am I aware of the problem of the time period and phase of networks of spiking excitatory and inhibitory spiking neurons having been posed or analysed before. Although there is a sizable literature on the dynamics of networks of spiking neurons, none of them appear to deal with networks comprising excitatory and inhibitory neurons [7,9]. While there has been much work on excitatory and inhibitory oscillators based on the population model of [11], this does not relate to spiking neurons, which allow different dynamics to continuous-valued population models [1]. By contrast, my extension to the locking theorem is trivial — indeed it has been hinted at [3]. Nevertheless, I believe that it has been a worthwhile accomplishment to have put it on a firmer footing.

Perhaps the most surprising result is the forbidden values of time period and phase, independent of the values of the weights. These regions possibly act to limit the time periods at which neural systems can work, and so are perhaps responsible for the time periods observed in the locust. The results also indicate that it should be rather difficult for the locust to maintain the experimentally-observed constant oscillation time period, even if the excitatory and inhibitory inputs are proportional to each other. This implies that either the input is relatively constant or that the oscillation time period does actually change in the locust, perhaps as a function of odour concentration. It may be possible that there are parameter regimes that do give the desired characteristics, however. It may also be possible that a more complex neuron model with active channels is required to achieve a fairly constant oscillation time period.

In deriving the results I have assumed that the refractory potential is strong enough to prevent PNs firing before they are inhibited. If this is not true, firing may not be "correct" as neurons may fire more than once on each cycle. Another, more limiting assumption is that of the homogeneity of the groups, a condition unlikely to be found in the nervous system. Less homogeneous groups will introduce differences in the firing times of units of a group, thus making the locking less precise. However, an extension of the one-group locking theorem to noisy neurons suggests that moderate amounts of noise still allow reasonably precise, stable locking and there are arguments as to why noisy neurons might be equivalent to inhomogeneous connections [1].

Several questions remain. In particular the effect of noise and inhomogeneity on the locking is not well understood. There may be ways of setting parameters to make large stable regions or to minimise the dependence of the time period on the inputs. The behaviour as the external inputs are increased in proportion to each other is relevant to the locust and should be studied.

In conclusion, this simple model of excitatory and inhibitory spiking neurons is a useful tool in predicting and understanding the activity of real neural systems without going to complex simulations. It may have application beyond the locust in some of the many excitatory-inhibitory subnetworks of nervous systems that show oscillatory activity.

A Proof of the Extended Locking Theorem

To show whether the group firing times are stable to small perturbations, we can perform a linear stability analysis. Suppose that the firing times before $\phi_l T$ are perturbed by small times $\delta_i^{(k)} \ll T \min_l \phi_l$. The firing times are now $(k+\phi_l)T + \delta_i^{(k)}$ for $k = 0, -1, -2, \ldots$ for a neuron in group $m < l$ and $(k+\phi_l)T + \delta_i^{(k)}$ for $k = -1, -2, \ldots$ for a neuron in group $m \geq l$.

The slight changes in firing times will change the membrane potential slightly. An increase in membrane potential when it is near the threshold makes a neuron fire earlier. The change in firing time is greater when the rate of change of the membrane potential is small. Mathematically,

$$\delta_i^{(0)} = -\frac{\delta h_i(\phi_l T)}{\dot{h}(\phi_l T)} \tag{19}$$

where $\delta h_i(\phi_l T)$ is the change in membrane potential due to the perturbations in the previous firing times.

We can linearise so that:

$$\eta_l(\phi_l T - (k+\phi_l)T - \delta_i^{(k)}) \approx \eta_l(-kT) - \dot{\eta}_{l,k}\delta_i^{(k)} \tag{20}$$

and

$$\epsilon_{lm}(\phi_l T - (k+\phi_m)T - \delta_i^{(k)}) \approx \epsilon_{lm}((\phi_l - \phi_m - k)T) - \dot{\epsilon}_{lm,k}\delta_i^{(k)} \tag{21}$$

where

$$\dot{\eta}_{l,k} = \dot{\eta}_l(-kT) \text{ and } \dot{\epsilon}_{lm,k} = \dot{\epsilon}_{lm}((\phi_l - \phi_m - k)T) \ .$$

To find the small change in membrane potential $\delta h_i(\phi_l T)$, we substitute the perturbed times for the unperturbed times in (6) and (7) and substitute the linearised expressions for the refractory, synaptic and membrane potential terms. Substituting the result in a rearranged version of (19) yields

$$\dot{h}_i(\phi_l T)\delta_i^{(0)} = \tag{22}$$
$$\sum_{k=-1}^{-F} \dot{\eta}_{l,k}\delta_i^{(k)} + \sum_{m=1}^{l-1}\sum_{j\in\mathcal{G}_m} J_{ij}\sum_{k=0}^{1-F}\dot{\epsilon}_{lm,k}\delta_j^{(k)} + \sum_{m=l}^{M}\sum_{j\in\mathcal{G}_m} J_{ij}\sum_{k=-1}^{-F}\dot{\epsilon}_{lm,k}\delta_j^{(k)} \ .$$

By the law of large numbers it is reasonable to assume that mean time shifts vanish, that is $(1/N)\sum_i \delta_i^{(k)} \approx 0$. In order to simplify the problem to make it tractable we must further assume that

$$\sum_{j\in\mathcal{G}_m} J_{ij}\delta_j^{(k)} \approx 0 \ . \tag{23}$$

Note that this is a stricter condition than (5) as it implies that the weights and perturbations are uncorrelated and that the weights sample the perturbations fairly well. For example, $J_{ij} = J_0\delta_{ij}$ (where δ is here the Dirac delta) satisfies (5) but does not necessarily lead to (23).

The assumption (23) allows us to neglect the last two summations in (22) so that the only delays influencing the 0th delay of neuron $i \in \mathcal{G}_l$ are the previous delays of neuron i itself:

$$\delta_i^{(0)} = \sum_{k=-1}^{-F} A_l^{(k)} \delta_i^{(k)} \text{ where } A_l^{(k)} = \frac{\dot{\eta}_{l,k}}{\sum_{k=-1}^{-F} \dot{\eta}_{l,k} + \dot{h}_i^{\text{syn}}(\phi_l T)} \ . \tag{24}$$

We can recast (24) in matrix form:

$$\begin{pmatrix} \delta_i^{(0)} \\ \vdots \\ \delta_i^{(1-F)} \end{pmatrix} = \mathbf{F}_l \begin{pmatrix} \delta_i^{(-1)} \\ \vdots \\ \delta_i^{(-F)} \end{pmatrix} \text{ where } \mathbf{F}_l = \begin{pmatrix} A_l^{(-1)} & A_l^{(-2)} & A_l^{(-3)} & \cdots & A_l^{(-F)} \\ 1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 0 \end{pmatrix} \ .$$

The condition for stability is that as $t \to \infty$, $\delta_i^{(k)} \to 0$. In the matrix notation, this is equivalent to $\lim_{n\to\infty} \mathbf{F}_l^n(\boldsymbol{\delta}) = 0$ for an arbitrary delay vector $\boldsymbol{\delta}$. A sufficient condition for this to happen is that the eigenvalues of $\mathbf{F}_l$ all lie within the unit circle. [3] use various matrix theorems to prove that $\sum_k A_l^{(k)} < 1$ is a necessary and sufficient condition for the modulus of the maximum eigenvalue to be less than one. Substituting (24) into this condition yields

$$\frac{\sum_{k=-1}^{-F} \dot{\eta}_{l,k}}{\sum_{k=-1}^{-F} \dot{\eta}_{l,k} + \dot{h}_i^{\text{syn}}(\phi_l T)} < 1 \ . \tag{25}$$

As, in the standard dynamics, $\dot{\eta}_{l,k} > 0$, the second part of the denominator must be positive, yielding the stability condition (11).

References

[1] Gerstner, W.: Population dynamics of spiking neurons: Fast transients, asynchronous states, and locking. Neural Computation **12** (2000) 43–89

[2] Gerstner, W., van Hemmen, J. L.: Associative memory in a network of 'spiking' neurons. Network: Computation in Neural Systems **3** (1992) 139–164

[3] Gerstner, W., van Hemmen, J. L., Cowan, J. D.: What matters in neuronal locking. Neural Computation **8** (1996) 1689–1712

[4] Laurent, G.: Dynamical representation of odors by oscillating and evolving neural assemblies. Trends in Neurosciences **19** (1996) 489–496

[5] Laurent, G., Davidowitz, H.: Encoding of olfactory information with oscillating neural assemblies. Science **265** (1994) 1872–1875

[6] Laurent, G., Naraghi, M.: Odorant-induced oscillations in the mushroom bodies of the locust. Journal of Neuroscience **14** (1994) 2993–3004

[7] Mirollo, R. E., Strogatz, S. H.: Synchronisation of pulse-coupled biological oscillators. SIAM Journal of Applied Mathematics **50** (1990) 1645–1662

[8] Smith, L. S., Cairns, D. E., Nischwitz, A.: Synchronization of integrate-and-fire neurons with delayed inhibitory connections. In *Proceedings of ICANN94*. Springer-Verlag (1994) pp. 142–145.
URL ftp://ftp.cs.stir.ac.uk/pub/staff/lss/icann94.ps.Z

[9] van Vreeswijk, C., Abbott, L. F., Ermentrout, G. B.: When inhibition not excitation synchronises neural firing. Journal of Computational Neuroscience **1** (1994) 313–321
[10] Wehr, M., Laurent, G.: Relationship between afferent and central temporal patterns in the locust olfactory system. Journal of Neuroscience **19** (1999) 381–390
[11] Wilson, H. R., Cowan, J. D.: Excitatory and inhibitory interactions in localized populations of model neurons. Biophysical Journal **12** (1972) 1–24